MoSt-DSA: Modeling Motion and Structural Interactions for Direct Multi-Frame Interpolation in DSA Images

Ziyang Xu^a, Huangxuan Zhao^b, Ziwei Cui^a, Wenyu Liu^a, Chuansheng Zheng^b and Xinggang Wang^{a,*}

^aInstitute of AI, School of EIC, Huazhong University of Science and Technology ^bUnion Hospital, Tongji Medical College, Huazhong University of Science and Technology

Abstract. Artificial intelligence has become a crucial tool for medical image analysis. As an advanced cerebral angiography technique, Digital Subtraction Angiography (DSA) poses a challenge where the radiation dose to humans is proportional to the image count. By reducing images and using AI interpolation instead, the radiation can be cut significantly. However, DSA images present more complex motion and structural features than natural scenes, making interpolation more challenging. We propose MoSt-DSA, the first work that uses deep learning for DSA frame interpolation. Unlike natural scene Video Frame Interpolation (VFI) methods that extract unclear or coarse-grained features, we devise a general module that models motion and structural context interactions between frames in an efficient full convolution manner by adjusting optimal context range and transforming contexts into linear functions. Benefiting from this, MoSt-DSA is also the first method that directly achieves any number of interpolations at any time steps with just one forward pass during both training and testing. We conduct extensive comparisons with 7 representative VFI models for interpolating 1 to 3 frames, MoSt-DSA demonstrates robust results across 470 DSA image sequences (each typically 152 images), with average SSIM over 0.93, average PSNR over 38 (standard deviations of less than 0.030 and 3.6, respectively), comprehensively achieving state-of-the-art performance in accuracy, speed, visual effect, and memory usage. Our code is available at https://github.com/ZyoungXu/MoSt-DSA.

1 Introduction

Frame interpolation, a class of fundamental tasks in computer vision, aims to deduce intermediate frames from given preceding and succeeding ones [11, 21]. These tasks are classified into single-frame and multi-frame interpolation based on the number of frames inferred [29, 14, 16]. Traditionally, multi-frame interpolation is achieved recursively. For instance, an intermediate frame I_b is inferred first, and then used with the ground truths of adjacent frames to deduce additional frames I_a and I_c [26, 23, 24]. However, this approach neither supports direct multi-frame interpolation nor allows flexible determination of frame count (typically odd).

DSA is an advanced medical imaging technology widely used in interventional surgery [25]. It is crucial for diagnosing and treating various vascular diseases, including brain, heart, and limbs. DSA operates by injecting a contrast agent, usually iodine-based, into the patient and capturing vascular images with X-rays. DSA technology varies: 2D DSA provides basic two-dimensional images. 3D DSA,

0.94 Memory Us MoSt-DSA FILM-L_{style} 2.20 G3.58 G 0.93 0.92 SoftSplat-EMA-VFI 0.91 0.9 0.1 0.2 0.3 0.4 0.5 0.6 0 Inference Time (s)

Figure 1. SSIM-Time-Memory comparison of different methods for direct interpolating 1 to 3 frames on our DSA dataset. Our

MoSt-DSA-L1 achieves 94.62, 94.35, 93.58 SSIM, 0.024s, 0.070s, 0.117s inference time, and 2.59G, 2.61G, 2.61G memory usage for interpolating 1 to 3 frames, respectively, outperforming SOTA EMA-VFI [38] in all aspects. Details in Tab. 1,2,3.

capturing images from multiple angles [39]. 4D DSA adds a time dimension, forming a sequence that captures dynamic blood flow's changes over time [12].

Moreover, frame interpolation for DSA images differs significantly from natural images. As comparing Fig. 2 with Fig. 3, DSA images present more complex structural and motion details [9]. Currently, no specific interpolation solutions for DSA images exist.

As we move from 2D to 4D DSA, frame interpolation complexity increases. Our research targets direct multi-frame interpolation for 4D DSA. Hereafter, DSA refers specifically to 4D DSA unless stated otherwise. Frame interpolation for DSA images confronts challenges from complex structures and motions. First, the vascular structure is complex: vessels are irregular, dense, and varied in size, like Fig. 3(a). Second, the imaging captures the contrast agent's diffusion, a non-rigid and complex motion depicted in Fig. 3(b). Third, vessels rotate during imaging, causing occlusions and overlaps that complicate motion analysis, as shown in Fig. 3(c).

To address the above challenges, extracting fine-grained and precise motion and structural features is critical. However, existing frame interpolation methods are tailored for natural scenes, resulting in unclear or coarse extraction of motion and structural features for



^{*} Corresponding Author. Email: xgwang@hust.edu.cn



Figure 2. Various motions in natural scenes. Motion subjects have simple structure and coarse texture feature granularity, also easy to predict the motion trajectory.





(b) Complex diffuse motion of contrast agents



(c) Complex rotational motion

Figure 3. Various challenges for frame interpolation in DSA images.

DSA images. Common approaches fall into three categories. The first uses a single module to mix and extract both motion and structural features, resulting in ambiguity in both aspects[16, 17, 21, 1, 5]. The second designs multiple modules to sequentially extract structural features of each frame and motion features between frames, although clear motion features are obtained, the corresponding structural relationships between frames are lacking[6, 40, 37, 4, 13, 22, 24, 26, 32, 36]. The third designs a single module to extract relative motion and structural features from frames simultaneously, but due to coarse context granularity, it fails to adapt to the fine-grained, complex structures of DSA images[38]. These methods commonly exhibit issues such as motion artifacts, structural dissipation, and blurring in DSA frame interpolation, as shown in Fig. 4.

In this work, we propose a network for flexible, efficient, and direct multi-frame interpolation in DSA images. Initially, we extract



Figure 4. Existing frame interpolation methods are tailored for natural scenes, and commonly exhibit issues such as motion artifacts, structural dissipation, and blurring in DSA frame interpolation.

multi-scale features from input frames and enhance them through cross-scale fusion. Inspired by the EMA-VFI [38], we propose a general module named MSFE that extracts motion and structural features between enhanced frames by cross-attention. Unlike EMA-VFI, MSFE doesn't rely on expensive attention maps and can flexibly adjust context-aware granularity. Specifically, by adjusting the optimal context range and transforming contexts into linear functions, MSFE calculates cross-attention between input frames in a fully convolutional manner, which reduces the storage cost and increases the computing speed. After extracting general motion and structural features through MSFE, we map the motion features at different times tand decode them together with the structural features to obtain flows and masks. Finally, a simplified UNet [28] refines features at different scales, decoding the flows, masks, and structural features to produce the corresponding intermediate frame I_t . A key advantage is that by extracting general motion and structural features only once, our MoSt-DSA can interpolate any intermediate frame by combining different time t during both training and testing. This is more flexible than methods interpolating for fixed t [24], more efficient than methods extracting different features for multiple t [16], and more direct than methods interpolating multi frames recursively [26, 23].

In summary, our work offers these main contributions:

- To our knowledge, MoSt-DSA is the first work that uses deep learning for DSA frame interpolation, and also the first method that directly achieves any number of interpolations at any time steps with just one forward pass during both training and testing.
- We propose a general module named MSFE that models motion and structural context interactions between frames by crossattention. Significantly, by adjusting the optimal context range and transforming contexts into linear functions, MSFE calculates cross-attention in a fully convolutional manner, which reduces the storage cost and increases the computing speed.
- We conduct extensive comparisons with 7 representative VFI models for interpolating 1 to 3 frames, MoSt-DSA demonstrates robust results across 470 DSA image sequences (each typically 152 images), with average SSIM over 0.93, average PSNR over 38 (standard deviations of less than 0.030 and 3.6, respectively), comprehensively achieving state-of-the-art performance in accuracy, speed, visual effect, and memory usage. If applied clinically, MoSt-DSA can significantly reduce the DSA radiation dose received by doctors and patients, lowering it by 50%, 67%, and 75% when interpolating 1 to 3 frames, respectively.

2 Related Work

2.1 Direct Multi-Frame Interpolation

Interpolation tasks for continuous image sequences, known as Video Frame Interpolation (VFI), aim to generate one or multiple intermediate frames between input frames [16, 14]. Considering that in DSA imaging, where the radiation dose correlates with image count, reducing frames and using AI interpolation instead can cut radiation significantly. Further, if multi-frame interpolation could be rapidly achieved with just one forward pass, it would not only further reduce radiation dose but also shorten the time consumed, securing more precious time for patient treatment. However, advanced VFI methods primarily focused on single-frame interpolation, with multiframe interpolation often reliant on recursion [38, 26, 23, 24]. During training, these methods are limited as they neither directly complete multi-frame interpolation nor allow flexible frame number determination, leading to a significant decrease in accuracy for direct multiframe interpolation during testing. Unlike these methods, our MoSt-DSA can directly achieve any number of frame interpolations at any time steps with just one forward pass during both training and testing.

2.2 Modeling Motion and Structural Interactions

Modeling the motion and structural interactions is essential for extracting motion and structural features. Existing frame interpolation methods are tailored for natural scenes, and the modeling of motion and structural interactions could divided into three categories. The first category concatenates input frames to a backbone network that extracts mixed features of motion and structure [16, 17, 21, 1, 5]. While straightforward to implement, these methods lack clear motion information, leading to restrictions in interpolating frames with various numbers and time steps [31, 8, 11]. The second category utilizes multiple modules to sequentially extract the structural features of each frame and the motion features between frames [6, 40, 37, 4, 13, 22, 24, 26, 32, 36]. Although these methods provide explicit motion features, they require modules with high computational costs, such as cost volume [13, 24, 26]. Moreover, capturing structural features from individual frames does not adequately identify the structural correspondence between frames, a critical aspect noted by [13] for VFI tasks. The third category, represented by [38], utilizes a single module for concurrent extraction of relative motion and structural features from frames. This approach's advantages include preserving and enhancing the detailed structural features of input frames without interference from motion features, mapping motion features to any moment for arbitrary intermediate frame generation, and significantly lowering training costs. However, due to coarse context granularity, it fails to adapt to the finegrained, complex structures of DSA images. These methods commonly exhibit issues such as motion artifacts, structural dissipation, and blurring in DSA frame interpolation, as shown in Fig. 4. Our method, aligning with the third category, introduces a general module named MSFE that models motion and structural context interactions between frames by cross-attention. Differing from [38], MSFE doesn't rely on expensive attention maps and can flexibly adjust context-aware granularity. By adjusting the optimal context range and transforming available contexts into linear functions, MSFE calculates cross-attention in a fully convolutional manner, which further reduces the storage cost and increases the computing speed.



Figure 5. An illustration of the MoSt Attention in the MSFE module for calculating motion and structural features. The enhanced structural features of I_0 and I_1 are involved in subsequent calculations in the MSFE module to generate the final structural features, see Fig. 7 for details.



Figure 6. An illustration of how we use Lambda Layer to calculate features in our MoSt Attention. Lambda Layer summarizes contextual information (within a scope r) into a fixed-size linear function (i.e. a matrix) applied to the corresponding query, thus bypassing the need for memory-intensive attention maps.

3 Method

Presenting a groundbreaking approach to direct multi-frame interpolation in DSA images, Fig. 7 delineates the overall network architecture of our method. Briefly, it is divided into five key modules. Initially, it employs the Multi-Scale Feature Extractor (**FE**), Cross-Scale Feature Fusion (**CSFF**), and Motion-Structure Feature Extractor (**MSFE**) to extract general motion and structural features. Subsequently, the Flow-Mask Estimator (**FME**) and **Refiner** decode and refine these features for different moments t to generate frame I_t .

3.1 Extracting General Motion and Structural Features

Multi-Scale Feature Extractor (FE). To excavate foundational features of blood vessels of various sizes before extracting motion and structural features, we first employ the FE to derive three different scales of neurovascular features. For input frames I_0 and I_1 , we initially compute the third layer of low-level features L_0^0 and L_1^0 , respectively, using 3x3 convolutions followed by PReLU [10]. Subsequently, through downsampling and the same convolution and activation configuration, we calculate the second layer of low-level features L_0^1 and L_1^1 , as well as the first layer of low-level features L_0^2 and L_1^2 for I_0 and I_1 respectively. Mathematically,



Figure 7. Overall network architecture of our MoSt-DSA. First, the Multi-Scale Feature Extractor (FE) processes the input frames I_0 and I_1 to obtain features at three different scales through continuous convolution and down-sampling. Next, Cross-Scale Feature Fusion (CSFF) uses multi-scale atrous convolutions [3] to generate fused features for I_0 and I_1 , which are linearly mapped and normalized. The Motion-Structure Feature Extractor (MSFE) then calculates motion and structural features from these fused features. Subsequently, for the intermediate time t, motion features are mapped. These mapped motion features, along with structural features and the original frames, feed into the Flow-Mask Estimator (FME) to predict flow and mask. Finally, the **Refiner** combines the various scale features from FE and structural features from MSFE, along with flow and mask, refining them into the image of intermediate time t.

$$\begin{cases} \boldsymbol{L}_{j}^{0} = \boldsymbol{H}\left(\boldsymbol{I}_{j}\right) \\ \boldsymbol{L}_{j}^{1} = \boldsymbol{D}\left(\boldsymbol{L}_{j}^{0}\right) \\ \boldsymbol{L}_{j}^{2} = \boldsymbol{D}\left(\boldsymbol{L}_{j}^{1}\right) \end{cases}, \qquad (1)$$

where H is a stack of convolution and activation functions, while D represents an integration of H with an additional downsampling operation, and j is 0 or 1.

Cross-Scale Feature Fusion (CSFF). To fuse neurovascular features of different scales and enhance the representation of foundational features, we further employ the CSFF for I_0 and I_1 to perform cross-scale feature fusion. Specifically, for the *i*-th layer lowlevel features L_0^i and L_1^i , we use 2^{i-1} atrous convolutions [3] (with a fixed kernel size of 3, stride of 2^i , and for the *n*-th atrous convolution, both padding and dilation size are *n*). Mathematically,

$$\boldsymbol{F}\left(\boldsymbol{L}_{j}^{i}\right) = \left(\boldsymbol{A}_{1}\left(\boldsymbol{L}_{j}^{i}\right), \dots, \boldsymbol{A}_{n}\left(\boldsymbol{L}_{j}^{i}\right)\right), \qquad (2)$$

where F signifies feature fusion, A indicates atrous convolution. The variable n, representing the number of atrous convolutions, takes a value of 2^{i-1} for i equal to 0, 1, or 2. Moreover, by merging the fused features from various scales and implementing a linear mapping, we obtain the cross-scale fused features F_0 and F_1 for I_0 and I_1 respectively, as:

$$\boldsymbol{F}_{j} = \mathcal{T}\left[\boldsymbol{\mathcal{C}}\left(\boldsymbol{F}\left(\boldsymbol{L}_{j}^{0}\right), \boldsymbol{F}\left(\boldsymbol{L}_{j}^{1}\right), \boldsymbol{F}\left(\boldsymbol{L}_{j}^{2}\right)\right)\right], \quad (3)$$

where \mathcal{T} represents the linear mapping, with \mathcal{C} indicates the concatenation operation. Finally, we flatten F_j and then normalize it, preparing for subsequent processing by the MSFE.

Motion-Structure Feature Extractor (MSFE). We propose MoSt Attention to calculate relative motion features while enhancing structural features between frames, as shown in Fig. 5 and 6 in detail. To facilitate and simplify understanding, the following formulas we give is based on our actual code implementation. We first concatenate F_0 and F_1 to obtain $F_a \in \mathbb{R}^{|n| \times d}$, and then acquire $F'_a \in \mathbb{R}^{|n| \times d}$ through reverse concatenation, as:

$$\begin{cases} \boldsymbol{F}_{a} = \boldsymbol{\mathcal{C}} \left(\boldsymbol{F}_{0}, \boldsymbol{F}_{1} \right) \\ \boldsymbol{F}_{a}^{\prime} = \boldsymbol{\mathcal{C}} \left(\boldsymbol{F}_{1}, \boldsymbol{F}_{0} \right) \end{cases}$$
(4)

Furthermore, we employ a Lambda Layer [2] to simulate contentbased and position-based contextual interactions in a fully convolutional manner. Specifically, we denote the depth of query and value as |k| and |v|, respectively, and denote the position information with $\boldsymbol{P} \in \mathbb{R}^{|n| \times d}$. The queries, keys, and values are calculated as follows:

$$\begin{cases} Q = F_a W_Q \in \mathbb{R}^{|n| \times |k|} \\ K = F'_a W_K \in \mathbb{R}^{|n| \times |k|} \\ V = F'_a W_V \in \mathbb{R}^{|n| \times |v|} \end{cases}$$
(5)

Then we represent relative position embeddings as $E \in \mathbb{R}^{|n| \times |k|}$. By normalizing the keys, we obtain $\overline{K} = \operatorname{softmax}(K, \operatorname{axis} = n)$. Next, we compute the content-based contextual interactions λ_c and position-based contextual interactions λ_p , as:

$$\begin{cases} \boldsymbol{\lambda}_{c} = \bar{\boldsymbol{K}}^{T} \boldsymbol{V} \in \mathbb{R}^{|k| \times |v|} \\ \boldsymbol{\lambda}_{p} = \boldsymbol{E}^{T} \boldsymbol{V} \in \mathbb{R}^{|k| \times |v|} \end{cases}$$
(6)

Finally, by applying contextual interactions to the queries as well as P, we obtain the general motion and structural features necessary for inferring any intermediate frame, as:

$$\begin{cases} S = Q\lambda_c + Q\lambda_p = \mathcal{C}(S_0, S_1) \\ M = P\lambda_c + P\lambda_p = \mathcal{C}(M_0, M_1) \end{cases} .$$
(7)

3.2 Decoding and Refining for Multi Intermediate Frames

To further obtain motion features corresponding to multiple different intermediate times t, we multiply each t with the general motion features to map and obtain $M_{0\to t}$ and $M_{1\to t}$.

Flow-Mask Estimator (FME). For a specific t, $M_{0\to t}$ and $M_{1\to t}$ are concatenated with S_0 and S_1 , respectively, and this combination serves as part a of the input for the FME, while part b is the concatenation of I_0 and I_1 . As shown in Fig. 7, FME (denote by \mathcal{F}) applies PixelShuffle [30] upsampling to part a, and downsampling to part b. Subsequently, parts a and b merge and undergo continuous convolution operations, eventually leading to the generation of bidirectional optical flow ϕ_t and mask μ_t corresponding to the specific t through upsampling, as:

$$\phi_t, \boldsymbol{\mu}_t = \mathcal{F}\left(\mathcal{C}\left(\boldsymbol{M}_{0 \to t}, \boldsymbol{M}_{1 \to t}, \boldsymbol{S}_0, \boldsymbol{S}_1\right), \mathcal{C}\left(\boldsymbol{I}_0, \boldsymbol{I}_1\right)\right). \quad (8)$$

Next, we initially employ ϕ_t to warp I_0 , I_1 , as well as the lowlevel features L_j^i from different layers extracted by FE, and the general structural features S_0 and S_1 extracted by MFSE. For instance, for X_z^y , the result after warping is denoted as $\widetilde{X_z^y}$. Subsequently, we concatenate I_0 , I_1 , \widetilde{I}_0 , \widetilde{I}_1 , ϕ_t , and μ_t together, referred to as \mathcal{O}_t .

Refiner. Finally, through the Refiner (a simplified UNet [28]), by integrating and refining features of different scales into O_t layer by layer, and then utilizing skip connection, we obtain the intermediate frame \hat{I}_t corresponding to t, as depicted in Fig. 7. Mathematically,

$$\widehat{I}_{t} = \widetilde{I}_{t} + R\left(\mathcal{O}_{t}, L, S\right), \qquad (9)$$

where R signifies the Refiner, L denotes the collection of L_{j}^{i} , and S refers to the collection of S_{0} and S_{1} . The symbol \odot represents the Hadamard product, and \tilde{t}_{i} is determined as follows:

$$I_{t} = \boldsymbol{\mu}_{t} \odot \text{backwarp} (I_{0}, \boldsymbol{\phi}_{t \to 0}) + (1 - \boldsymbol{\mu}_{t}) \odot \text{backwarp} (I_{1}, \boldsymbol{\phi}_{t \to 1}).$$
(10)

3.3 Loss Functions

To further enhance the inference quality, we employed a combination of three types of loss functions, as follows:

$$\mathcal{L} = \boldsymbol{w}_{1}\mathcal{L}_{1} + \boldsymbol{w}_{\text{VGG}}\mathcal{L}_{\text{VGG}} + \boldsymbol{w}_{\text{Style}}\mathcal{L}_{\text{Style}}, \quad (11)$$

where \mathcal{L}_1 denotes the L1 reconstruction loss, which minimizes the pixel-wise RGB difference. Additionally, \mathcal{L}_{VGG} employs the L1 norm of the VGG-19 features to enhance finer image details and texture quality [33]. The style loss \mathcal{L}_{Style} utilizes the L2 norm of the auto-correlation of the VGG-19 features [7, 27, 18]. This approach aims to further leverage the benefits of \mathcal{L}_{VGG} by capturing and replicating style patterns and textures more effectively. Regarding the selection of the weights (w_1 , w_{VGG} , w_{Style}), we referenced [26].

4 Experiments

4.1 Datasets

We collected 470 head DSA image sequences from 8 hospitals, each from a different patient, typically containing 152 images of 489x489 resolution. These were split into 329 for training and 141 for testing, maintaining a 7:3 ratio. For each sequence targeting *n*-frame interpolation, we arrange it into several groups, each with consecutive n + 2 frames. Adjacent groups start one frame apart. For details regarding data acquisition, we use NeuAngio33C, NeuAngio43C, and NeuAngio-CT equipment, following the SpinDSA protocol.

4.2 Implementation Details

Model Configuration. For interpolating 1 to 3 frames, time (t) sequences are set to [0.5], [0.33, 0.67], and [0.25, 0.50, 0.75], respectively. For simulating contextual interactions, context modeling scope (r) sizes are 29, 29, and 21. Effects of varying r are compared in the ablation study.

Training Details. We trained on 4 A100 GPUs, and for tasks interpolating 1 to 3 frames, we set the batch sizes to 10, 10, and 6, with warm-up steps set to 9000, 12000, and 16000, respectively. We use the AdamW [20] optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a weight decay of 1e - 4. The learning rate is warmed up to 2e - 4 and then decays following a cosine schedule [19], decreasing to 2e - 5over 300 epochs. We crop each frame to a resolution of 320×320 and apply random flip and rotation for augmentation. Regarding the selection of loss weights (w_1 , w_{VGG} , w_{Style}), we referenced [26], assigning weights of (1.0, 1.0, 0.0) for the first epoch and weights of (1.0, 0.25, 40.0) for the subsequent epochs.

Testing Details. To highlight our method's advantages, we compared MoSt-DSA with representative VFI methods. For ABME [24] and SoftSplat [23], we tested on released pre-trained weights due to the absence of training codes. For EMA-VFI (state-of-the-art) [38] and FILM [26], we retrained them on our dataset following their original setups. All tests were performed on a single RTX 3090 GPU.

Comparison Details. We trained two versions: one (MoSt-DSA- \mathcal{L}_1) using only the \mathcal{L}_1 loss, which achieves higher test scores; the other (MoSt-DSA) using our proposed combined loss \mathcal{L} , which benefits image quality (see supplementary materials [35] for proof). When comparing visual effects, we use the version of the model that yields high image quality[26, 7], i.e., FILM- \mathcal{L}_{Style} , and SoftSplat- \mathcal{L}_F .



Figure 8. Visual comparison for interpolating one frame with methods in VFL. " \star " indicates the SOTA method in VFI. The first and third rows (on the right side) correspond to the green box in blend, and the second and fourth rows (on the right side) correspond to the blue box in blend. M stands for motion artifact, S for structural dissipation, and B for blurring. By comparing the inference and residual results of various methods, it can be proved that the natural scene VFI method has many problems of motion artifact, structural dissipation and blurring, while MoSt-DSA relatively obviously alleviates these problems.

Table 1.	Quantitative comparison with VFI methods on single-frame
ir	terpolation. Best scores for color losses in blue, and for
manaamtu	ally consistive losses in red. The second lowest memory users in

perceptually-sensitive losses in **red**. The second lowest memory usage in **green**. "*" indicates the SOTA VFI method. EMA is short for EMA-VFI.

	SSIM (%)		PSNR		Time (s)	Memory
	Mean↑	STD↓	Mean↑	STD↓	1 frame↓	1 frame↓
ABME [24]	94.02	2.28	39.83	3.39	0.383	2.66 G
FILM- \mathcal{L}_1 [26]	94.11	2.37	39.86	3.43	0.201	3.51 G
EMA-small [38]	94.19	2.21	40.07	3.47	0.027	2.10 G
EMA* [38]	94.33	2.13	40.13	3.40	0.056	2.63 G
MoSt-DSA- \mathcal{L}_1	94.62	2.12	40.32	3.35	0.024	2.59 G
SoftSplat- \mathcal{L}_F [23]	91.78	3.07	38.59	3.51	0.035	2.20 G
FILM- \mathcal{L}_{VGG} [26]	93.10	2.67	39.27	3.45	0.201	3.51 G
FILM- \mathcal{L}_{Style} [26]	93.05	2.72	39.25	3.48	0.201	3.51 G
MoSt-DSA	93.65	2.61	39.55	3.45	0.024	2.59 G

 Table 2. Quantitative comparison with VFI methods on two frames interpolation. The meanings of blue, red, green, and "*" are the same as those in Table 1. EMA is short for EMA-VFI.

	SSIM (%)		PSNR		Time (s)	Memory
	Mean↑	$\text{STD}{\downarrow}$	Mean↑	$\text{STD}{\downarrow}$	2 frame \downarrow	2 frame↓
FILM-L ₁ [26]	92.62	3.13	37.93	3.82	0.388	3.58 G
EMA-small [38]	91.82	3.68	37.37	4.07	0.074	2.10 G
EMA* [38]	91.90	3.63	37.41	4.08	0.112	2.64 G
MoSt-DSA- \mathcal{L}_1	94.35	2.29	39.78	3.44	0.070	2.61 G
SoftSplat- \mathcal{L}_F [23]	90.86	3.49	37.84	3.43	0.084	2.20 G
FILM- \mathcal{L}_{VGG} [26]	91.42	3.42	37.47	3.67	0.388	3.58 G
FILM- \mathcal{L}_{Style} [26]	91.31	3.50	37.39	3.74	0.388	3.58 G
MoSt-DSA	93.14	2.84	38.94	3.38	0.070	2.61 G

4.3 Single-Frame Interpolation

We visualized the single-frame interpolation results of each model and compared them with the ground truth by calculating residuals.

As shown in Fig. 8, the first and third rows correspond to the green box in blend, and the second and fourth rows correspond to the blue box in blend. M stands for motion artifact, S for structural dissipa-

Table 3. Quantitative comparison with VFI methods on three frames interpolation. The meanings of blue, red, green, and "*" are the same as those in Table 1. EMA is short for EMA-VFI.

	SSIM (%)		PSNR		Time (s)	Memory
	Mean↑	STD↓	Mean↑	STD↓	3 frame↓	3 frame↓
FILM-L ₁ [26]	91.94	3.52	37.21	3.91	0.548	3.58 G
EMA-small [38]	90.48	4.52	36.31	4.24	0.122	2.10 G
EMA* [38]	90.57	4.50	36.35	4.24	0.165	2.64 G
$MoSt-DSA-\mathcal{L}_1$	93.58	2.69	38.85	3.56	0.117	2.61 G
SoftSplat- \mathcal{L}_F [23]	90.07	3.87	37.24	3.53	0.137	2.20 G
FILM- \mathcal{L}_{VGG} [26]	90.63	3.83	36.75	3.74	0.548	3.58 G
FILM- \mathcal{L}_{Style} [26]	90.54	3.90	36.66	3.82	0.548	3.58 G
MoSt-DSA	93.03	2.94	38.66	3.59	0.117	2.61 G



Figure 9. Intuitive comparison of metrics for each method interpolating 1 to 3 frames. Our MoSt-DSA- \mathcal{L}_1 leads SOTA EMA-VFI in all respects, while showing superior robustness with a lower STD.

tion, and B for blurring. By comparing the results of various methods, it can be proved that the natural scene VFI method has many problems of motion artifact, structural dissipation and blurring, while MoSt-DSA relatively obviously alleviates these problems.

The quantitative comparison for single-frame interpolation, as shown in Tab. 1, demonstrates our MoSt-DSA's superiority in SSIM,



Figure 10. Impact of context modeling scope (r) sizes for DSA frame interpolation tasks: from 1 to 3 frames. The first to third columns correspond to interpolating 1 to 3 frames, and the first to second rows represent the mean and STD, respectively. The stable STD proves the robustness of MSFE, and the Mean indicates that the best r for interpolating frames 1 to 3 is 29, 29, and 21, respectively.

PSNR, and inference time over all competitors. Furthermore, our model also boasts more efficient memory usage than EMA-VFI during inference. Notably, the score differences among SOTA methods in the VFI domain are minimal. For instance, on the UCF101 dataset [34], the top-performing EMA-VFI surpasses the second-best [15] by only 0.01% in SSIM and 0.01 in PSNR, and the third-best [40] by 0.04% in SSIM and 0.01 in PSNR. Thus, it is a significant margin that our MoSt-DSA- \mathcal{L}_1 's lead over the EMA-VFI, by **0.29%** in SSIM and **0.19** in PSNR, as shown in Tab. 1.

4.4 Direct Multi-Frame Interpolation

We further compared our method with representative VFI methods in tasks of direct interpolating 2 and 3 frames.

For each method, we set t=[0.33, 0.67] and t=[0.25, 0.50, 0.75] for interpolating 2 to 3 frames, respectively. Considering that ABME [24] couldn't interpolate at arbitrary time steps, we excluded it from the comparison. We give the average values across metrics for each frame count. For instance, if interpolating 2 frames results in SSIM values of [0.8, 0.9], then the average is 0.85.

The quantitative evaluation results for interpolating 2 and 3 frames are presented in Tab. 2 and 3. Our MoSt-DSA continues to outperform other methods, in terms of SSIM, PSNR, and inference time, also exhibiting a lower standard deviation (STD). Memory usage during inference also remains more efficient than the EMA-VFI. This conclusively demonstrates the superior robustness of our method.

We further intuitively compared the metrics for interpolating 1 to 3 frames, Fig. 9 shows a clear pattern: MoSt-DSA's superiority in SSIM, PSNR, and stability grows with the increase in interpolated frames. Compared to EMA-VFI, MoSt-DSA- \mathcal{L}_1 's SSIM is higher by **0.29%**, **2.45%**, and **3.01%** for interpolating 1, 2, and 3 frames. In terms of PSNR, the increase is **0.19**, **2.37**, and **2.50**. Furthermore, MoSt-DSA- \mathcal{L}_1 's STD for SSIM is lower by 0.4%, 37%, and 40%, and for PSNR, it is 1.49%, 16%, and 16% lower. We believe this significant lead reflects the advantages of MoSt-DSA trained with

multi-frame supervision to model motion and structural interactions accurately, and highlights the importance of multi-frame supervision training for direct multi-frame interpolation tasks.

4.5 3D Reconstruction Showcase from Single Frame Interpolation

We conducted 3D reconstructions using both the single-frame interpolated sequences (interpolating every other image) and the original sequences. Our results are virtually indistinguishable to the reconstruction from original data. Details in supplementary materials [35].

4.6 Ablation Study

Impact of context modeling scope (r). Fig. 10 demonstrates that the impact of r on STD is minimal, highlighting MSFE's robustness. r's influence is slightly more pronounced on Mean-of-SSIM (no more than 1.8%) than on Mean-of-PSNR (no more than 1.1). More detailed numerical results are available in the supplementary materials [35].

Loss function comparison on our MoSt-DSA. We prove that our proposed loss function significantly improves image quality, in the supplementary materials [35].

5 Conclusion

We have proposed MoSt-DSA, the first work that uses deep learning for DSA frame interpolation, to reduce radiation dose in DSA imaging significantly. In particular, we devised a general module that models motion and structural context interactions between frames in a fully convolutional manner, by adjusting the optimal context range and transforming available contexts into linear functions. Experiment results show that our MoSt-DSA outperforms the state-ofthe-art Video Frame Interpolation methods in accuracy, speed, visual effect, and memory usage for interpolating 1 to 3 frames, and can also assist physicians in 3D diagnosis and treatment.

Acknowledgements

This work is supported by the National Key R&D Program of China (No. 2023YFC2413500).

References

- W. Bao, W.-S. Lai, C. Ma, X. Zhang, Z. Gao, and M.-H. Yang. Depthaware video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3703–3712, 2019.
- [2] I. Bello. Lambdanetworks: Modeling long-range interactions without attention. *arXiv preprint arXiv:2102.08602*, 2021.
- [3] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [4] D. Danier, F. Zhang, and D. Bull. St-mfnet: A spatio-temporal multiflow network for frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3521– 3531, 2022.
- [5] T. Ding, L. Liang, Z. Zhu, and I. Zharkov. Cdfi: Compressiondriven network design for frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8001–8011, 2021.
- [6] P. Gao, H. Tian, and J. Qin. Video frame interpolation with flow transformer. arXiv preprint arXiv:2307.16144, 2023.
- [7] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference* on computer vision and pattern recognition, pages 2414–2423, 2016.
- [8] S. Gui, C. Wang, Q. Chen, and D. Tao. Featureflow: Robust video interpolation via structure-to-texture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14004–14013, 2020.
- [9] Y. Guo, L. Bi, E. Ahn, D. Feng, Q. Wang, and J. Kim. A spatiotemporal volumetric interpolation network for 4d dynamic medical image. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020.
- [10] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), December 2015.
- [11] Z. Huang, T. Zhang, W. Heng, B. Shi, and S. Zhou. Rife: Real-time intermediate flow estimation for video frame interpolation. arXiv preprint arXiv:2011.06294, 2020.
- [12] S. Ito, M. Kanagaki, N. Yoshimoto, Y. Hijikata, M. Shimizu, and H. Kimura. Cerebral proliferative angiopathy depicted by fourdimensional computed tomographic angiography: A case report. *Radiology Case Reports*, 2022.
- [13] Z. Jia, Y. Lu, and H. Li. Neighbor correspondence matching for flowbased video frame synthesis. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5389–5397, 2022.
- [14] H. Jiang, D. Sun, V. Jampani, M.-H. Yang, E. Learned-Miller, and J. Kautz. Super slomo: High quality estimation of multiple intermediate frames for video interpolation. In *Proceedings of the IEEE conference* on computer vision and pattern recognition, pages 9000–9008, 2018.
- [15] X. Jin, L. Wu, J. Chen, Y. Chen, J. Koo, and C.-h. Hahm. A unified pyramid recurrent network for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1578–1587, 2023.
- [16] T. Kalluri, D. Pathak, M. Chandraker, and D. Tran. Flavr: Flow-agnostic video representations for fast frame interpolation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (WACV), pages 2071–2082, January 2023.
- [17] L. Kong, B. Jiang, D. Luo, W. Chu, X. Huang, Y. Tai, C. Wang, and J. Yang. Ifrnet: Intermediate feature refine network for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2022.
- [18] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*, pages 85–100, 2018.
- [19] I. Loshchilov and F. Hutter. Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983, 2016.
- [20] I. Loshchilov and F. Hutter. Decoupled weight decay regularization, 2019. URL https://arxiv.org/abs/1711.05101.

- [21] L. Lu, R. Wu, H. Lin, J. Lu, and J. Jia. Video frame interpolation with transformer. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3532–3542, 2022.
- [22] S. Niklaus and F. Liu. Context-aware synthesis for video frame interpolation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1701–1710, 2018.
- [23] S. Niklaus and F. Liu. Softmax splatting for video frame interpolation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5437–5446, 2020.
- [24] J. Park, C. Lee, and C.-S. Kim. Asymmetric bilateral motion estimation for video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14539–14548, 2021.
- [25] A. Posa, A. Tanzilli, P. Barbieri, L. Steri, F. Arbia, G. Mazza, V. Longo, and R. Iezzi. Digital subtraction angiography (dsa) technical and diagnostic aspects in the study of lower limb arteries. *Radiation*, 2022.
- [26] F. Reda, J. Kontkanen, E. Tabellion, D. Sun, C. Pantofaru, and B. Curless. Film: Frame interpolation for large motion. arXiv preprint arXiv:2202.04901, 2022.
- [27] F. A. Reda, G. Liu, K. J. Shih, R. Kirby, J. Barker, D. Tarjan, A. Tao, and B. Catanzaro. Sdc-net: Video prediction using spatially-displaced convolution. In *Proceedings of the European conference on computer* vision (ECCV), pages 718–733, 2018.
- [28] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference* on Medical image computing and computer-assisted intervention, pages 234–241. Springer, 2015.
- [29] W. Shang, D. Ren, Y. Yang, H. Zhang, K. Ma, and W. Zuo. Joint video multi-frame interpolation and deblurring under unknown exposure time. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 13935–13944, June 2023.
- [30] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time single image and video superresolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.
- [31] Z. Shi, X. Xu, X. Liu, J. Chen, and M.-H. Yang. Video frame interpolation transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17482–17491, 2022.
- [32] H. Sim, J. Oh, and M. Kim. Xvfi: Extreme video frame interpolation. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 14489–14498, 2021.
- [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- [34] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402, 2012.
- [35] Z. Xu, H. Zhao, Z. Cui, W. Liu, C. Zheng, and X. Wang. Most-dsa: Modeling motion and structural interactions for direct multi-frame interpolation in dsa images, 2024. URL https://arxiv.org/abs/2407.07078.
- [36] T. Xue, B. Chen, J. Wu, D. Wei, and W. T. Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 2019.
- [37] Z. Yu, Y. Zhang, D. Zou, X. Chen, J. S. Ren, and S. Ren. Rangenullspace video frame interpolation with focalized motion estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22159–22168, 2023.
- [38] G. Zhang, Y. Zhu, H. Wang, Y. Chen, G. Wu, and L. Wang. Extracting motion and appearance via inter-frame attention for efficient video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5682–5692, 2023.
- [39] H. Zhao, Z. Zhou, F. Wu, D. Xiang, H. Zhao, W. Zhang, L. Li, Z. Li, J. Huang, H. Hu, et al. Self-supervised learning enables 3d digital subtraction angiography reconstruction from ultra-sparse 2d projection views: a multicenter study. *Cell Reports Medicine*, 2022.
- [40] C. Zhou, J. Liu, J. Tang, and G. Wu. Video frame interpolation with densely queried bilateral correlation. arXiv preprint arXiv:2304.13596, 2023.